

Method and System for Sending Information on an Extranet

Inventor

Michael E. Gaddis

Prepared by:

Fenwick & West LLP
Two Palo Alto Square
Palo Alto, CA 94306
Attorney Docket No. 22013-04959

Express Mail No.: EL482473371US

17

Method and System for Sending Information on an Extranet

9/11

Background

A. Technical Field

This application relates to a method of transferring data over multiple networks and, more specifically, to a method and pricing scheme for brokering data sent and received over an extranet that provides a financial incentive for network providers to carry each other's customer traffic in a quality fashion.

B. Background of the Invention

10 i. The Need for High Quality Internet Service

In certain situations, it is desirable to send data over a network with a high priority and with a guaranteed maximum transit time. For example, certain data may be needed in real time or may be of high importance. Currently, certain conventional network protocols (such as the Asynchronous Transfer Mode (ATM) protocol) contain provisions for indicating a "level of service" that particular transmitted data is to receive -- a capability referred to as Quality of Service (QoS). Users pay premiums to obtain higher levels of service in an ATM network. It would be desirable to send data over the Internet with the same type of guarantees. Unfortunately, the design of the Internet does not provide financial incentives for Internet Service Providers (ISPs) to cooperate in a way that would result in performance guarantees for the users of the Internet.

15
20
ii. Barriers to the Deployment of Hardware Solutions

One possible way to accomplish this goal of high quality service on the Internet would be to upgrade the routers used to route Internet traffic. Unfortunately, the deployment of Quality of Service (QoS) capable routers end-to-end in the Internet would require a massive investment. Making such a radical upgrade is not currently practical even if the carriers were motivated to do so. This creates a classic "chicken and egg" problem as to which will come first -- the investment for the QoS network upgrades or the acceptance of a QoS service and the incremental revenue to pay for that investment.

30 The barriers to the deployment of QoS, therefore, are currently substantial. First, deployment of QoS requires a massive investment in network infrastructure. Second, there are

currently no exchange services to facilitate the transfer between ISPs even if two ISPs made that investment. Third, the lack of current economic drivers (financial incentives to the ISPs) makes the necessary investment highly risky for ISPs or potential clients.

5 iii. The Peering Conundrum

10 The exchange of traffic between Internet service providers (ISPs), where traffic from one ISP's customers is destined for customers of another ISP, is referred to as "peering." This is contrasted with "transit," where the traffic that is exchanged is destined not only for the receiving ISPs customers, but also other users throughout the Internet (which are in turn reached via the receiving ISPs peering connections). ISPs generally charge for transit connections, while peering connections do not involve an exchange of money. So, peering provides free bi-directional exchange of customer traffic between ISPs whereas transit provides paid access to the entire Internet. For example, large ISPs with their own national backbones often mutually agree with other large ISPs with national backbones to freely exchange their customer data (Note: An Internet backbone is the central part of an ISPs network which transports data between edge, or regional, parts of the networks and Internet peering points). Instead of peering, large ISPs that have a national backbone offer transit service to smaller, regional ISPs who cannot offer reciprocal backbone sharing. These smaller ISPs use the larger ISPs national backbones to reach users outside of their service areas. Peering connections are based upon the underlying assumption that the traffic flows between two different ISPs will be approximately equal. These mutual agreements to exchange information traffic freely and without charge do not make economic sense if one peering partner is forced to carry substantially more traffic than another. Unbalanced bandwidth may affect peering relationships because most peered data is routed between peers using a routing method called "hot potato routing."

25 In conventional Internet peering, an originating peer network carrying a customer's traffic finds the nearest entry point to the destination peer's network and drops the traffic off as soon as possible. Under this so-called "hot potato" routing, it is advantageous for the originating carrier to "get rid of" traffic as soon as possible by passing it to a peer. There is no incentive for an ISP to transport additional bandwidth if a peer network is available to carry the traffic. Likewise, the destination peer network will send returning data to the originator's network from its nearest entry point into the originator's network. In an ideal world, If the two

exchange points used for this data transfer are geographically dispersed, this type of routing creates an asymmetric routing path between two network terminations whereby one “to” path travels primarily on the destination peer’s network and the “from” path travels primarily on the originator’s network. Ideally, the amount of bandwidth used is symmetric in both directions and the two networks have similar geographic scope, resulting in a fair economic exchange. Each peer carries half the bandwidth and each peer has (presumably) one paying customer to pay for its half of the transfer.

Unfortunately, although Internet peering models are generally based upon the assumption of symmetric bandwidth between peers, this assumption has proven to be flawed. The growth of the World Wide web (WWW) and increases in WWW traffic have exacerbated the problem of bandwidth asymmetry. Web content providers tend to transmit copious amounts of data. Bandwidth between Web content providers and the individuals who view Web content is often unbalanced. Web content providers typically send as much as 4 to 10 times more bandwidth than they receive. An individual might request a Web page by sending just a single address or page request, whereas the Web page content provider returns large amounts of Web page data to the individual, thereby consuming large amounts of bandwidth.

For example, consider a scenario between ISP A, servicing an individual, and ISP B, servicing a Web content provider. ISP A and B both use hot-potato routing, dumping network traffic off at the closest entry point to the peering ISP. A client of ISP A sends a request for a Web page, and ISP A routes the request such that the request is carried by ISP B for a majority of the distance traveled (hot potato routing). ISP B returns the requested Web page, constituting a considerably larger amount of bandwidth than the initial request from the individual. ISP B similarly uses hot potato routing, which results in ISP A carrying the Web traffic for the majority of the distance traveled. ISP A is forced to carry more than its fair share of the bandwidth traffic burden in this scenario.

In the above example, ISP A is unable to charge its own customers more for the extra bandwidth, because the entity originating the extra bandwidth is a customer of ISP B. Thus, the current peering system does not provide the proper economic incentives for an ISP to increase its bandwidth or its quality of service, because the increased bandwidth and quality may be consumed by the ISP’s peers.

The above-described situation often results in unacceptably poor service for business-to-business (B2B) traffic on the Internet. Many business customers have declined to migrate their strategic network systems from Frame Relay, ATM and private lines to the much more cost-effective public IP Internet because the internet cannot provide the performance and service guarantees they require. This lack of willingness to send data via the Internet has slowed the acceptance of Internet Virtual Private Networks (VPNs). If this lack of quality and confidence is left unchecked, it will slow Internet market segment growth into the B-2-B commerce market, which is estimated to exceed \$7.3 trillion in B2B e-commerce transactions in the coming years.

What is needed is a system and method that motivate ISPs to provide the QoS required by their Internet Service Provider (ISP) customers who support B2B and other types of premium data delivery.

Summary of Embodiments of the Invention

The described embodiments of the present invention offer motivation and incentive for a network provider, such as an ISP, to provide guaranteed Quality of Service (QoS) to its clients, whether those clients are individual users, other carriers, or companies.

The potential exists even in today's network to completely revolutionize B2B communications using the Public IP Internet if certain performance requirements can be met. Business customers need assurances from Internet carriers in the form of end-to-end Service Level Agreements (SLAs) that have teeth and that cover multiple carriers. Guarantees for availability, packet loss, delay, and throughput must be met. When these requirements are met, a majority of the traffic currently in private B2B networks can migrate to the public IP Internet. Once this occurs, Internet carriers will generate more incremental revenue and business customers are able to consolidate communications and reduce cost.

The described embodiments of the present invention are based on a belief that a significant number of business customers will pay an additional fee for end-to-end guarantees between networked business destinations because that incremental cost is far less than they pay in private network infrastructure or in lost business due to poor Internet performance. At least one described embodiment of the present invention is further based on a belief that a premium Public IP service may bypass traditional peering connectivity, replacing it with a different

pricing and delivery model. In certain embodiments, a premium public IP service may be offered in addition to regular Internet service to the same customers over the same physical infrastructure.

The described embodiments of the present invention implement monetary payments based on exchange points in a backbone structure to build incentives for quality exchange between customers of the backbone carrier. It is envisioned that, even if the described embodiments of the invention do not totally prevail, networks of the future will have more than one type of exchange capability. For example, there might exist a capability for best effort services (peering), one for premium services (using bulk monetary payments) and in the long run, a system for per flow payments for real-time performance guarantees for voice and video (a subset of the overall flow in the network).

The described embodiments of the present invention are based on a belief that an ISP must be paid for all premium class data entering its network in order to build incentives for that carrier to carry the data to the remote transit end of its network. This arrangement requires transmitting carriers to pay a payment to the receiving network and the receiving network to meet certain performance requirements to warrant that payment. The transmitting carrier will facilitate that exchange between the sending network and the receiving network, verify the quality and determine the fees between carriers and charge a small brokerage fee for its trouble.

Advantages of the invention will be set forth in part in the description which follows and in part will be apparent from the description or may be learned by practice of the invention. The objects and advantages of the invention will be realized and attained by means of the elements and combinations particularly pointed out in the appended claims and equivalents.

Brief Description of the Drawings

Fig. 1 is a block diagram showing a conventional Internet end-user model.

Fig. 2 is a block diagram showing a conventional Internet transit model.

Fig. 3 is a block diagram showing a conventional Internet peering model.

Fig. 4 is a block diagram showing a conventional Internet Multi-Transit/Peering model.

Fig. 5 is a block diagram showing an Incentive model in accordance with a preferred embodiment of the present invention.

Fig. 6(a) is a block diagram showing an Internet Data Exchange System having multiple customers in accordance with a preferred embodiment of the present invention.

Fig. 6(b) is a block diagram showing additional details of a part of the Internet Data Exchange System of Fig. 6(a).

Fig. 7(a) is a block diagram indicating a flow of funds in between several ISPs and a backbone carrier using an Incentive method implementation in accordance with the present invention.

Fig. 7(b) is a block diagram indicating a flow of funds in between several ISPs and a backbone carrier in accordance with a combination of an Incentive method implementation the present invention and with Internet peering.

Fig. 8(a) is a block diagram showing a central collection point collecting data from routers connected to a backbone carrier in accordance with the present invention.

Fig. 8(b) is a block diagram showing system performance measurements being made.

Figs. 9(a) and 9(e) are flow charts showing a method of gathering data and determining payments in accordance with the data.

Figs. 9(b)-9(d) provide an example of data gathered to determine a 95th percentile for data originated and the data terminated for a given customer.

Fig. 9(f) is a flow chart showing a method of determining a performance for various system elements.

Fig. 9(g) is an example of the type of performance statistics gathered for the purposes of testing conformance with an SLA.

Fig. 9(h) is an example of a table of required performance values.

Detailed Description of Embodiments

Reference will now be made in detail to several embodiments of the present invention, examples of which are illustrated in the accompanying drawings. Wherever practicable, the same reference numbers will be used throughout the drawings to refer to the same or like parts.

i. Internet Routing

Figs. 1-4 demonstrate functionality available in conventional Internet systems. Fig. 1 is a block diagram showing a conventional Internet end-user model. In this conventional model, end users, such as businesses 110, 120 purchase access to the Internet from regional or national

Internet Service Providers (ISPs) 130, paying a monthly fee to their ISP for the amount of bandwidth they use (for example, a usage-based, or "burstable" connection) or expect to use (i.e., a maximum amount of capacity is set by the ISP). This bandwidth is generally offered as a partial or full T1 or T3 circuit, measured in terms of the megabits per second (Mbps) carried or available in either direction. Note that, if the end users 110 and 120 both use the same ISP, they usually do not connect to the Internet at all. Data can be sent bi-directionally. That is, data can be sent and received simultaneously.

Fig. 2 is a block diagram showing a conventional Internet transit model. In this model, an end user 210 connects to a regional ISP 230 and an end user 220 connects to a national ISP 240. Regional ISPs, such as ISP 230 must purchase Internet access from one or more of the national ISPs 240 in order to carry their customers' data to most destinations, since the network of a regional ISP does not have the scope to connect to them directly. This type of Internet access is called a "transit" connection, since the regional ISPs 230 transit the network of the national ISPs 240 to get to the destination end-user 220. Regional ISPs 230 purchase transit Internet access from larger ISPs much the same way that business end-users do; they simply pay a monthly fee for amount of bandwidth they use or expect to use. These transit connections between a regional and a national ISP are typically partial or full T3 circuits, and are measured in terms of the Mbps in either direction. Thus, in the transit model, a regional ISP pays a fee to a national ISP that covers both incoming and outgoing traffic of the regional ISP.

In the situation shown in Fig. 2, regional ISP 230 bears the following costs:

- Edge "A" cost: cost of transmitting data within the regional ISP 230 network.
- Transit cost for connecting to the national ISP's network

In the situation shown in Fig. 2, national ISP 240 bears the following costs:

- Long-haul cost: cost of transmitting data within the national ISP 240 network.
- Edge "B" costs: cost of transmitting data to the end user 220 within the national ISP

240 network. Thus, in this model, user 210 pays IS P 230 to send and receive data and ISP 230 pays ISP 240 to send and receive data.

Fig. 3 is a block diagram showing a conventional Internet peering model. In this model, an end user 310 connects to an originating ISP (such as a regional ISP 330) and an end user 320 connects to another originating ISP (such as regional ISP 342). In this example, ISPs 330 and 342 act as both originating and terminating ISPs, since their customers both send and receive

data, although this may not always be the case for all ISPs. In terms of the flow of funds from the customers, the system of Fig. 3 is essentially a “bill and keep system” whereby each ISP bills its respective end users for data sent and received and keeps the revenues. The lack of economic incentives for carrying each other’s traffic presents an obstacle to offering end-to-end

5 performance guarantees, since each ISP generally tries to “get rid of” data to a peer as quickly as possible to minimize the costs of carrying the traffic. Due to the way the Internet Protocol (IP) works, neither the originating nor the terminating ISP knows exactly where a packet of data is destined on the other provider’s network. Therefore, ISPs generally take a data packet to the closest “peering point” and transfer it to the destination ISPs network and that ISP carries it on
10 to the destination end-user (i.e., their customer). This is referred to as “hot potato” routing, because it effectively gets traffic off of the originating ISP’s network as quickly as possible. Once traffic is off-net of the originating ISP, the originating ISP cannot assure performance.

In Fig. 3, when data is sent from user 310 to user 320, it is passed from originating ISP 330 to ISP 342, which acts as a “long-haul” ISP, transporting the data to end user 320. This
5 occurs because peering ISPs generally hand off data to a peer ISP as quickly as possible. Similarly, when data is sent from user 320 to user 310, it is passed from originating ISP 342 to ISP 330, which acts as a “long-haul” ISP, transporting the data to end user 310. Again, ISP 342 generally uses hot potato routing.

As shown, most national ISPs currently exchange traffic using a “peering” connection, in
20 which neither party pays the other party for the connection. It is assumed that each party will send/receive an equal amount of traffic to/from the other (i.e., they are “peers”) such that if they were to charge each other similar fees, the balance for each would be zero. For this reason, large/national ISPs usually do not peer with small/regional ISPs. Instead, they require regional ISPs to pay for access, as shown in Fig. 2. Even so, potential inequities exist in peering
25 arrangements between national ISPs because push/pull traffic is usually not balanced. For example, users tend to download much more data than they upload. As a somewhat simplified example, if one ISP supports a web server and its peer ISP supports multiple users viewing the web, the data traffic will most likely be unequally distributed, with one ISP sending much more data than it is receiving.

30 In Fig. 3, each originating ISP bears the cost of:

–Edge transport on originating traffic: transporting packets before handing off to their peer ISP.

–Edge plus long-haul transport on terminating traffic: transporting packets after they are handed off from an originating ISP.

Fig. 4 is a block diagram showing a conventional Internet Multi-Transit/Peering model.

This model involves some combination of transit and peering. In Fig. 4, when data is sent from user 410 to user 420, it is passed from regional ISP 430 to regional ISP 440, and then to national ISP 450, which acts as a “long-haul” ISP, transporting the data to end user 420. A Multi-Transit/Peering model has two or more ISPs at either or both edges of a data transmission. In this example, the regional ISPs use a transit model to pass data, while regional ISP 440 has a peering relationship with national ISP 450. In another example, regional ISPs 430 and 440 might be peered, while regional ISP 440 has a transit relationship with national ISP 450. In the example, any of ISPs 430, 440, or 450 can be the long-haul ISP.

Thus, conventional Internet traffic uses one of 1) a transit model, 2) a peering model, or 3) a combination of both, as depicted in the Multi-Transit/Peering model of Fig. 4.

i. An Incentive Model

It will be understood that certain diagrams herein are somewhat simplified for clarity for explanation. For example, many systems in accordance with the invention include multiple exchanges, multiple customers per exchanges, and multiple clients per customer.

Fig. 5 is a block diagram showing an Incentive model in accordance with a preferred embodiment of the present invention. In this model, end users 510 and 520 still contract with their local ISPs for service as described above. The end users pay their ISPs for access and the ISPs carry traffic to and from the end users. ISP 530 and ISP 550 are called “customers” of the backbone carrier 540. The relationships between ISP 530 and backbone carrier 540 and between ISP 550 and backbone carrier 540, however, are not based (at least entirely) on conventional peering relationships.

Instead, under the Incentive model, each customer of backbone 540 is debited (charged) for originating traffic and credited (paid) for terminating traffic to and from the backbone. These debits/credits are based on the amount of data transmitted, as measured in specific Megabits per

second (Mbps) increments. Thus, for example, an ISP might promise to pay backbone carrier 540 a fixed amount to transmit 1 Mbps. In certain embodiments, there is an understanding between the customer and backbone carrier 540 that the data would be transferred at least at a predetermined rate (e.g., a predetermined Mbps rate) in accordance with an SLA, as described below.

In the Incentive model, a part of the amount debited from the originating customer is used to credit the terminating customer. The remainder of the part of the amount debited from the originating customer constitutes a fee that backbone carrier 540 receives for brokering the transaction between originating customer 530 and terminating customer 550. In Fig. 5, an originating customer is defined as a customer, such as an ISP, that places the data on the backbone 540. In Fig. 4, a terminating customer is defined as a customer, such as an ISP, that receives data from backbone 540. It will be understood that there can be other networks involved in data transmission. For example, as shown in Fig. 5, data might transverse more than one ISP in a transit and/or peering model before being transferred by the originating ISP to backbone 540. Similarly, data might transverse more than one ISP before being transferred to an end user after the data leaves backbone 540. Thus, in Fig. 5, the terms originating customer and terminating customer are defined as customers that “touch” backbone 540 and, respectively, deliver data to and from the backbone. Customers are not limited to only ISPs and can be, for example, ISPs, companies, or individuals.

In certain embodiments in accordance with the system of Fig. 5, backbone 540 agrees to a Service Level Agreement (SLA) with each customer. Backbone 540 withholds credit for terminating traffic if the terminating ISP does not meet the agreed-upon SLA. In certain embodiments, the withheld credits are then transferred to the originating ISPs as compensation. In certain embodiments, backbone carrier 540 also agrees to certain performance requirements and if backbone carrier 540 does not meet its requirements, it is not credited with its fee for brokering the transaction. If the backbone carrier 540 does not meet its requirements, the backbone carrier still credits the terminating customer/ISP. The backbone carrier’s brokerage fee is generally collected, but may be returned to the originating customer or the terminating customer (or split between the originating and terminating customers). Such refund may be made, for example, on a monthly basis. The SLAs between the various customers and backbone 540 may differ from each other, since they preferably are negotiated separately between

backbone 540 and the various customers. Backbone 540 may have different SLAs with the various ISPs. It is contemplated that the SLAs will generally be similar and it is also possible to have a universal SLA. Even so, it is not always required that all SLAs be the same. For example, certain large customers may be able to negotiate more favorable SLAs with the owners of the backbone carrier 540.

In another embodiment, certain customers are involved in "portfolio SLAs" with the backbone carrier. In a portfolio SLA, performance is measured for a group of one or more clients, not for all clients as described above. Thus, for example, two large companies who exchange large amounts of data between themselves might desire higher performance for that traffic. In this situation, performance data is collected specifically for traffic between these two companies. Debits and credits in the Incentive model for those companies in the SLA portfolio are made based on measures of performance between the companies in the SLA portfolio.

The described Incentive model creates a framework for implementing inter-carrier Quality of Service (QoS), whereby different classes of service may be tailored to specific applications (voice, video, etc.) and with different SLAs offered by the participating ISPs. The debits/credits are different for these classes of service, as are (presumably) the prices charged to the end-users. Alternatively, a backbone carrier might implement an Incentive model having a single class of service. This single class of service would preferably be better than the conventional Internet service. Alternately, a backbone might implement an Incentive model having tiered charges for different levels of usage (such as different volumes of traffic). Thus, a user who transmitting or receive a large amount of traffic would pay a lower rate (but a higher total amount) than a user who transmitted a small amount of traffic. The amounts paid to the terminating customer in this situation might stay the same or might also vary in accordance with the tier of the originating customer. Debits or credits also may be based on exact usage or on tiers of usage. Alternately, the backbone carrier may implement an Incentive model having an unlimited pricing structure in which all customers are debited a flat fee (for example, on a monthly basis). In such a model, it might also occur that all terminating customers are paid a credited fee (for example, on a monthly basis). Alternately, the backbone carrier may decide not to charge certain large customers or clients for a predetermined time period in order to convince those large customers or clients of the value of the QoS offered by the Incentive model. Thus,

for example, a large potential client might not be charged for traffic during a period of time. An ISP/customer might also be convinced not to charge the potential client during the same period.

In the Incentive model of Fig. 5, backbone 540 incurs all transport costs, including long-haul transport costs, manages credits/debits to customers, and administers end-to-end SLAs to ensure off-net quality. Each customer/ISP 510, 550 bears the edge cost for their respective customers.

The Incentive model of Fig. 5 is preferably embodied by an "Internet Data Exchange System" (iDES)SM. Fig. 6(a) is a block diagram showing an Internet Data Exchange System having a backbone 540 and multiple customers 530, 550 in accordance with a preferred embodiment of the present invention. A typical iDES preferably includes the following components:

- NADs – Network Analysis Devices

- High speed backbone routers (ERs) located at Extranet Exchange Sites (EE), such as ER 610 and ER 612, and

- Backbone 540 connected to the routers (ER) 610, 612. The term dark fiber backbone refers to infrastructure that is in place but not yet in use. In general, backbone carriers that own or have an Indefeasible Right of Use (IRU) for a fiber optic network use the term "dark fiber" to denote that they originally lit the fiber and have complete operational control over its use. For example, some electric utilities have installed optical fiber cable where they already have power lines installed in the expectation that they can lease the infrastructure to telephone or cable TV companies or use it to interconnect their own offices. In the described embodiment backbone 540 includes more than 10,000 route miles of dark fiber and uses Nortel Optera lightwave technology to light this fiber, although other physical implementations of backbone 540 may be used without departing from the spirit and scope of the invention.

In Fig 6(a), backbone structure 540 has at least two exchange sites 610, 612. Customers, such as regional ISPs 530, 550, connect to backbone structure 540 via these exchange sites. In the example, clients A and C connect to exchange sites of backbone 540 via regional networks 530 and 550 (customers) by establishing respective tunnels 602 and 604 on regional networks 530 and 550. In this embodiment, the tunnel is preferably implemented using the GRE tunneling protocol as defined in at least RFC 1701, RFC 1702, and RFC 2784 (GRE over IP), which are herein incorporated by reference. "RFCs" ("Request for Comments") are documents available

from the Internet Engineering Taskforce (IETF) defining various aspects of Internet design and operation. Any appropriate protocol may be used to implement tunneling, including but not limited to Microsoft's PPTP protocol or Cisco System's Layer 2 Forwarding protocol. Use of a tunnel between a client and an exchange site helps ensure that the customer's data is delivered to the backbone carrier. Similarly, use of a network Analysis Device (NAD) allows for measurement of the network performance between a client and an Exchange site. A NAD can be used to perform other services than performance measurement, such as encryption and/or Voiceover IP, which may affect the debits/credits/ and brokerage fees of the Incentive model, since encryption and/or voice-over IP can affect system performance.

In the example, client B, who requests and sends data over the network, connects to customer 530 without a tunnel. In the example, client B also does not include an NAD. Client B is a client of customer 530. Client B is included in this example to illustrate that not all connections to an exchange point need to be made via a tunnel. Client B might also include a NAD. Any client in the system might include a NAD.

In the example, client D, who requests and sends data over the network, connects directly to router 610 without a tunnel. Client D is included in this example to illustrate that some clients, such as, for example, corporations and certain intranets, connect to an exchange point without tunneling and without going through an ISP/customer network. In this case, client D can also be considered a customer. Each of customers 530 and 550 has at least one client. Most clients have a NAD, which collects performance data as discussed below. Fig. 6(a) also includes various routers (R) within the Internet transit system and various routers (C) within backbone 540. Fig. 6(a) also includes NADs for active monitoring of network performance.

Fig. 6(b) is a block diagram showing additional details of an example part of the Internet Data Exchange System of Fig. 6(a). Two customer ISPs 530, 550 are connected to backbone 540 via respective exchange points 610, 612. In the embodiment shown, the connection between the customers 530, 550 and the exchange points 610, 612 are a Synchronous Optical Network (SONET) over optical fiber. The base rate (OC-1) is 51.84 Mbps. OC-2 runs at twice the base rate, OC-3 at three times the base rate, and so forth. The system of Fig. 6(b) also includes NADs to monitor and measure performance and Network Management Systems, to implement SLAs, as described below.

Fig. 7(a) is a block diagram indicating an example of flow of funds between two customers (e.g., ISPs 710 and 720) and a backbone 740 in accordance with an implementation of an Incentive method of the present invention. Each customer/ISP has one or more clients, as shown.

5 In the example of Fig. 7(a), ISP A 710 has contracted with the carrier who owns backbone 740 for the debit and credit amounts paid, respectively, for originating and terminating data from and by ISP A. In this example, ISP A is debited/charged \$475 for every Mb of data that it originates to backbone 740. Similarly, ISP A is credited/paid \$190 for every Mb of data that it terminates from backbone 740. Importantly, ISP A does not negotiate with ISP B.

10 Instead, each customer/ISP negotiates with backbone 740 to agree upon a debit/credit scheme. In at least one embodiment, all customers have the same debit/credit scheme with the backbone, but this is not a requirement for all embodiments.

Thus, in the example of Fig. 7(a), ISP A charges its client#1 \$1900 for every 1Mb that client#1 sends via ISP A. In the example, ISP A has an agreement with client#1 that ISP A will deliver certain premium data rates and reliability measures in exchange for this \$1900 fee. When client#1 sends 1Mb of data to ISP A 710, ISP A 710 passes the data to backbone 740, which transmits it as shown, for example, in Fig. 5 and passes it to the destination ISP, such as ISP B 720. In this example, backbone 740 debits ISP A \$475 when it receives the originating data from ISP A. Backbone 740 credits ISP B 720 with \$190 when the 1Mb of data is delivered to ISP B. In certain embodiments, this credit is paid only if ISP B 720 meets certain performance measurements as specified in the SLA between backbone 740 and ISP B. Backbone carrier 740 retains \$285 of the money debited from ISP A for itself as a brokerage fee ($\$475 - \$190 = \$285$).

In this example, ISPs A and B have identical debit/credit arrangements with backbone 740. Thus, in this example, ISP B charges its client#2 \$1900 for every 1Mb that client#2 sends via ISP B. ISP B also has an agreement with client#2 that ISP B will deliver certain premium data rates and reliability measures in exchange for this fee. When client#2 sends 1Mb of data to ISP B 720, ISP B 710 passes the data to backbone 740, which transmits it as shown, for example, in Fig. 5 and passes it to the destination ISP, such as ISP A 710. In this example, backbone 740 debits ISP B \$475 when it receives the originating data from ISP B. Assuming that ISP A 710 meets certain performance measurements as specified in the SLA between backbone 740 and ISP A, backbone 740 credits ISP a 710 with \$190 when the 1Mb of data is delivered to ISP A.

Backbone 740 credits ISP A 710 with \$190 when the 1Mb of data is delivered to ISP A. In certain embodiments, this credit is paid only if ISP A 710 meets certain performance measurements as specified in the SLA between backbone 740 and ISP A. Backbone carrier 740 retains \$285 of the money debited from ISP A for itself as a brokerage fee ($\$475 - \$190 = \$285$).

Thus, as shown in Fig. 7(a), of the \$1900 paid by client#1, when data is sent, \$1425 is retained by the originating ISP ($\$1900 - \475 paid to the backbone = \$1425). In the example, this is true for both ISP A and ISP B. Of the \$1900 paid by client#1, \$475 is initially debited from ISP A. Of this \$475, \$285 is retained by backbone 740 as a brokerage fee and \$190 is credited to the terminating customer/ISP. Thus, in the example, if ISP A both sends and receives 1Mb of data via backbone 740, ISP A will receive a total of \$1615 for sending 1Mb of data and receiving 1Mb of data ($\$1900$ (from client#1) + \$190 (for terminating data) - \$475 (debited for originating data) = \$1615). In the example, ISP B receives the same sums. In the example, when either of ISP A and ISP B is the originating ISP, backbone 740 retains a brokerage fee of \$285 for each Mb of data it carries.

It should be understood that the example of Fig. 7(a) is not to be taken in a limiting sense. In other implementations, backbone 740 may have a different fee and SLA arrangements with the two ISPs and may adjust its brokerage fee accordingly. In both of these examples, the sum of the amount retained by backbone 740 plus the sum credited to the receiving ISP is equal to the amount debited from the originating ISP. In other implementations, however, these three values (debit, credit, and brokerage fee) need not add up as in the example. For example, backbone 740 might decide to operate at a small loss for a period of time. Similarly, backbone 740 might reserve a part of its brokerage fee to pass on to the ISPs as a reward or incentive or it might reserve a part of its brokerage fee to pay to some third party as a reward or incentive. In general, in the described embodiments of the Incentive method, the amount debited from the originating customer/ISP is larger than the amount credited to the terminating customer/ISP.

As discussed above, use of an Incentive method affords several advantages. These advantages include the ability of backbone 740 to offer a much-needed QoS solution to ISPs/customers and thus, to increase the satisfaction and retention of those ISPs/customers and their clients. Moreover, the incremental revenue from offering premium priced service and additional bandwidth, hardware and services to clients benefit all ISPs involved in the system. Increased QoS tends to help ISPs retain clients and the ISPs can profit more from the premium

service, even though they lose some of the premium service fee to the brokerage fee for backbone 740. Moreover, the described embodiments reduce long-haul transport costs to customer/ISPs. Moreover, it is contemplated that the customers/ISPs will benefit from the marketing and sales support from the owner of backbone 540 carrier.

5 The described Incentive method also provides a cost effective solution for mission critical data communications with end-to-end performance guarantees. The fact that, in the implementation shown, the Incentive method is implemented on the Internet provides a ubiquitous access and any-to-any connectivity for the clients of the ISPs that are the customers of the backbone 740. It should be understood that the Incentive model is not necessarily limited to
10 the Internet, but instead, could be used on any appropriate network, whether wired or wireless.

Fig. 7(b) is a block diagram indicating a flow of funds in between several ISPs and a backbone carrier in accordance with an Incentive method. In the example, the ISPs have also allowed their clients to sign up for regular, non-premium Internet service, which does not involve an Incentive model. In this example, the Incentive model works as described above in connection with Figure 7(a). In the example, when data is sent via the traditional Internet model, instead of collecting \$1900 for premium service, ISP A also collects a (presumably) smaller amount, such as \$1400, for 1Mb of Internet access through conventional Internet peering. As described above, in Internet peering, ISP A and ISP B exchange data to and from each other without charging each other for the data transmission. While such a peering arrangement eventually delivers the data, it does not provide the same performance guarantees as the Incentive method. Thus, in this example, when it transmits 1Mb of data via an Incentive method and 1Mb of data via the conventional internet model, ISP A retains the full \$1400 paid by its client#1 for conventional access plus \$1615 as described in connection with Fig. 7(a). Note that
20 this example assumes that client#1 will pay higher rates to ISP A to ensure a higher standard a
25 service and reliability under the Incentive method. As discussed above, the agreements between backbone 740 and ISP A and between backbone 740 and ISP B can be, but need not be the same.

iii. The Effect of SLAs on the Incentive Model

Although the Incentive model can be practiced without using SLAs, SLAs add incentive
30 for the customers (and the backbone) to maintain high levels of service, since they will not be paid if their level of service drops below that specified in the relevant SLA. In a system that

does not use SLAs, debits and credits are made from and to originating and terminating customers without regard to performance measurements.

Fig. 8(a) is a block diagram showing a central collection point 850 collecting data from routers of a backbone carrier in accordance with the present invention. In the described embodiment, the routers in the system of Fig. 8 preferably are Juniper Networks model M20 routers, manufactured by Juniper Networks of Mountain View, CA, although other appropriate routers and network devices may be used without departing from the spirit and scope of the invention. Routers used in the described embodiment are able to periodically determine their throughput (Megabits per second, i.e., Mbps) and communicate the determined throughput to a central location, such as a data center or main server. For example, in the described embodiment, the routers 862, 864, 866 at the exchange sites of the backbone have the capability of determining an average throughput every five minutes for data sent and for data received. These average throughput determinations are then transmitted to a central site 850. Specifically, in the described embodiment, each router periodically transmits a number of bytes and a number of packets sent and received during a current time period. The number of bytes and packets are determined for each tunnel connection between the ISP and an exchange point. In the described embodiment, the SNMP protocol is used to transport the data to the central point 850.

Fig. 8(b) is a block diagram showing system performance measurements being made. In the example, clients have NADs to monitor and measure performance of test messages that are periodically sent through the system to test the performance of the system overall (i.e., end-to-end) and of the individual customers and the backbone. The performance data includes information identifying where the test messages originate from, how long it takes them to pass through the originating and terminating customers and the backbone (latency), and what packet loss and jitter are associated with the transmission. Jitter is described as the amount that the transmission rate actually varies from the mean during a current time period. In the example, shown, the average time it takes a test message to pass through customer/ISP A is X1 milliseconds; to pass through backbone 740 is Y1 milliseconds, and to pass through customer/ISP B is Z1 milliseconds. Similarly, respective customer/ISP A, backbone 740, and customer/ISP B have respective packet losses during the measured period of X2%, Y2%, and Z2%. Similarly, the availability of customer/ISP A, backbone 740, and customer/ISP B is X3%, Y3%, and Z3%, respectively. Similarly, the throughput of customer/ISP A, backbone 740, and

customer/ISP B is X4Mb, Y4Mb, and Z4Mb, respectively. Similarly, the jitter of customer/ISP A, backbone 740, and customer/ISP B is X5 millisecs, Y5 millisecs, and Z5 millisecs, respectively.

This information is preferably used to determine if particular elements of the system have met their SLA requirements. Note that, in the described embodiments, performance measurements are made separately from measurements of transmitted data. Other embodiments may measure performance by making measurements on the transmitted data. In the described embodiment, performance data is measure for latency, packet loss, and availability. These statistics are kept for the backbone, for each customer, and for the system overall (i.e., end-to-end).

Figs. 9(a) and 9(e) are flow charts showing a method of gathering data and determining payments in accordance with the data. In Fig. 9(a), during a current billing period, such as a month, a central location receives data regarding a number of originated and terminated bits from the exchanges/routers. This data indicates, for example, the client by whom the traffic was sent or received, and the amount of traffic. The customer associated with each client is known (or can be determined). This data is added to a database as shown in Fig. 9(b). In Fig 9(a), when the billing period is over, the central location determines a 95th percentile of originating bits and a 95th percentile of terminating bits for each customer/ISP. As seen below, the 95th percentile is used to determine the amounts to debit and credit each customer.

Fig. 9(b) is an example of data gathered to determine a 95th percentile for data originated and the data terminated for a given customer. The figure shows a table containing data for a single client of a single customer during a single billing period. It will be understood that similar data is gathered for each client. For each client, a number of originated Mbits and a number of terminated Mbits are stored for each sample period in the billing period. At the end of the billing period, this data is used to determine a 95th percentile of originated data and terminated data for that client as shown in Figs. 9(c) and 9(d).

Fig. 9(c) shows that the sample originating data for a client are sorted in descending order. The top 5th percentile of sorted samples is ignored and the next largest sample -- the 95th percentile sample -- is used to calculate the debit/charge (originating) for the client. The 95th percentile (originating) for all clients of a customer is combined to form the 95th percentile (originating) for the customer. As shown in Fig. 9(d), a similar process is used to determine the

95th percentile (terminating) for all clients of a customer and to combine these values to form the 95th percentile (terminating) for the customer. As seen below, the 95th percentiles (originating and terminating) of the customer are used to determine the amounts to debit and credit the customer.

5 There are three important factors to a percentile calculation:

1) The percentile number

10 A percentile basically says, that, for that percentage of the time, the data points are below the resulting value. A 95th percentile says that 95% of the time data points are below that value and 5% of the time they are above that value. In a system planning for the mean use or average use, the network could become unusable (saturated) half the time, therefore the described embodiment plans for a higher usage rate (95%).

2) Data points used

15 A percentile is calculated on some set of data points. What those data points represent is significant to understanding the meaning of the percentile result. Network percentiles are based on sampled throughput utilization. The sample rate indicates how accurate or forgiving the percentile is. The more frequent the sample rate, the more accurate and less forgiving the percentile will be. In the described embodiment, data samples are collected every five minutes (a sample period). The routers count bits over a 20 5 minute period, and the data sample represents a five minute averaged bits per second value. Because the value received from the routers is averaged, the highs and lows within that 5 minute period are not known.

3) Data set size

25 The data set size indicates the range of the values. In network percentiles, the data set is a period of time over which samples are collected. Usually for any solid planning and trend determination, we need a reasonably large data set to cover the peaks and valleys of utilization.

30 In the described embodiment, the exchanges send samples every 5 minutes and the billing period is approximately one month, although any reasonable sample rate and billing period can

be used. Thus, the number of samples from the exchange/routers for each client can be approximately:

$$12 \times 24 \text{ (hours)} \times 30 \text{ (days)} = 8640 \text{ samples}$$

Thus, a sample from an exchange/router includes information about the number of Mb originated and terminated to/from that exchange/router for a tunnel of each customer in the last 5 minutes. 8640 of these samples (taken over the billing period) are used to determine the 95th percentile of a client, and the clients' 95th percentiles are used to determine the 95th percentile of customers.

At least one embodiment of the present invention uses "tiered" pricing. If a customer originates (or terminates) between 0 and a first number of bytes, the customer is charged a first amount per byte. If the customer originates (or terminates) above the first number of bytes during the billing period, but less than a second, higher number, the customer is charged a different amount per byte. In the described embodiment, the prices per byte decrease for higher tiers. Other variations of tiered pricing are possible and can be used with the Incentive model discussed herein. For example, a different tiered scheme could be used for terminating and originating data. Alternately, other embodiments might charge a flat fee per byte, without using tiered pricing.

Fig. 9(e) is a flow chart showing a method of determining payment amounts (i.e., debit and credit amounts) for each customer. Element 951 begins a loop for each exchange site. Element 952 begins a loop for each tunnel through a customer. (The connections of the client B to an exchange and of the combined customer/client D to an exchange of Fig. 6(a) preferably are counted as a tunnel for these purposes). In element 953, for the current exchange site and the current tunnel to/from a client, if the current customer originated data, the current customer is debited in accordance with the 95th percentile of data that the customer originated during the billing period. (In another embodiment, the customer could be debited in accordance with his actual originated number of bytes and the 95th percentile calculation is not used.)

Element 954 determines, if the customer terminated data for this tunnel, whether the terminating customer met performance standards specified in its SLA with the backbone. If the terminating customer did not meet the SLA during the sample period, then in element 958, the backbone carrier preferably retains the entire amount that would have been credited to the terminating customer. (In certain embodiments, part of this amount is shared with the

originating customer as compensation if the terminating customer's performance also fails to live up to the SLA of the originating customer with the SLA). If the terminating customer met the SLA during the sample period, then in element 956, the backbone carrier credits the customer for data that it terminated during the sample time period in accordance with the 95th percentile for this customer's terminated data. Thus, the customer would be credited by an amount yielded by multiplying the 95th percentile of terminated data for that customer by a unit price.

Elements 960-970 relate to circumstances in which the backbone carrier may not retain its brokerage fee. It is contemplated that, in the described embodiment, the backbone always collects a brokerage fee, although it might not retain it. Element 960 determines whether the backbone met its individual performance SLA with the originating customer. If not, the brokerage fee that would normally be retained by the backbone is given to the originating customer in element 966 (or in some embodiments, split between the originating and terminating customers as compensation). In certain circumstances, the customers are encouraged to pass this amount on to their clients, although this is not a requirement.

If, in element 964, the backbone met its SLA, the backbone retains a brokerage fee that is part of the amount debited from the originating customer. Elements 968 and 970 end their respective loops.

Fig. 9(f) is a flow chart showing a method of determining performance for various system elements. In element 980, the average latency for the backbone and for each customer is determined. In the described embodiment, latency is measured in milliseconds and represents a time for a roundtrip from one client to another. As discussed above, the clients regularly send test packets and transmit the resulting measurements to the central location for a determination of average latency.

In element 982, the average packet loss for the backbone and for each customer is determined. In the described embodiment, packet loss is measured in percentages of packets sent. As discussed above, the clients regularly send test packets and transmit the resulting measurements to the central location for a determination of average packet loss.

In element 984, the percentage of availability for the backbone and for each customer is determined. In the described embodiment, availability is measured in percentages of total time. As discussed above, the clients regularly send test packets and transmit the resulting measurements to the central location for a determination of average availability.

In element 986, the average latency, packet loss, and availability for the overall system is determined. This number is obtained by combining the performance results for individual parts of the system. At least element 986 is optional.

Fig. 9(g) is an example of the type of performance statistics gathered for the purposes of testing conformance with an SLA. The figure shows that average latency, packet loss, and availability are determined for the backbone and for each customer and for each client. Fig. 9(h) is an example of a table of required performance values. In the described embodiment, all customers have the same SLA with the backbone, although this is not always the case. Thus, for all SLAs, the maximum acceptable average latency for the backbone is 70ms; the maximum acceptable packet loss for the backbone is .2%; and the minimum acceptable availability for the backbone is 99.99%. Similarly, under all SLAs, the maximum acceptable average latency for a customer is 40ms; the maximum acceptable packet loss for a customer is .4%; and the minimum acceptable availability for a customer is 99.9%. Note that other embodiments may have different acceptable values, depending on the SLAs negotiated between the customers and the backbone carrier. In the figure, the maximum acceptable average latency for the overall system is 150ms; the maximum acceptable packet loss for the overall system is 1%; and the minimum acceptable availability for the overall system is 99.8%.

As shown in elements 954 and 958 of Fig. 9(c) if a terminating customers fail to meet their SLA performance requirements, is does not receive credit for terminating data. Thus, terminating customers are motivated to delivery high quality service. As shown in element 964 of Fig. 9(c), if the backbone fails to meet its SLA performance requirements, the backbone may not retain its brokerage fee. Thus, the backbone carrier is motivated to delivery high quality service and not to retain customers that hurt the quality of service on the overall system.

In the described embodiment, a typical SLA between the backbone carrier and ISP/customer would have minimum requirements for one or more of: maximum time without availability of the network; packet loss; throughput; bit error rate; latency; and jitter. Not all SLAs necessarily have requirements for each of these values. Other SLAs might have different measures of performance.

As discussed above, the existence of premium service class(es) will enable the development of the QoS Internet by providing the economic incentive to invest in next generation routers and switches to provide QoS. It will also give rise to unique service

constructions such as direct ISP connections for premium consumer dial-in, dedicated DSL for SOHO (small office, home office) or telecommuting and other uses as technology evolves.

For example, online brokerages might desire high quality service to offer high quality dedicated dial-in accounts for their day-traders and platinum customers. They could achieve this by partnering with a national dial-in that attaches to backbone 740 with a high level SLA. With these mechanisms in place they could then establish a guaranteed stock trading performance window (which is their ultimate goal).

Because of the way in which the service is constructed, with the backbone cut-through with carrier edge access, the deployment of QoS switching equipment will likely occur at the edges of the Internet in an incremental manner. This will mitigate the cost burden for deploying the QoS Internet while creating the economic drivers to encourage its growth.

iv. Exchange Sites

Exchange sites are constructed through the interconnection of many ISPs at a regional point-of-presence of backbone 740. These connections might be high capacity local loops like DS-3s and OC-3s or another type of multi-megabit pipe; or, it may via a part/extension of the backbone's dark fiber network. With these basic interconnections providing physical connectivity, routers are connected from the Internet carriers to an exchange site router.

In the premium services category, a Border Gateway protocol (BGP) may be used to announce some, all, or none of the IP addresses that are part of the global premium services group (all addresses of customers that have purchased one of these services). This will be determined depending on the method of cut-through interconnect used to connect a regionally attached premium services customer to an assigned regional exchange site. For instance, if tunnels are used to facilitate forced regional routing then only the end-points of the tunnels need be announced and the customer addresses that are reachable over that tunnel would only be known by the backbone carrier and the specific customer. However, if routing preferences were used by the Internet carrier to effect regional routing then a full BGP announcement might be necessary. Various embodiments may implement routing in any appropriate way without departing from the spirit and scope of the invention.

v. Testing and Verification

An important part of the present invention is the ability to measure and determine bandwidth usage and performance. SLAs require the ability to test performance and report the results in a meaningful way.

Certain embodiments of the invention include an advanced SLA testing service for premium service customers. This will be accomplished by deploying inexpensive servers, referred herein as "network analysis devices (NADs)," at the client's site to run tests between itself, NADs within the exchange site and to other NADs at other client sites.

In addition to the performance measurement method described above, it is presumed that clients and customers might use the standard tools like the Unix ping command (for availability and latency) and trace command (to insure that regional routing is effective). Other implementations will also test transport protocols like TCP and UDP to insure effective performance exists at those levels as well and will develop the capability to test higher level protocol planes like WEB, video and voice flows.

Thus, in summary, the current invention includes a payment method that allows a backbone carrier to broker data between originating and terminating points, where the originating point is debited and the terminating point is credited, and where the backbone owner retains some part of the difference between the debited amount and the credited amount. Furthermore, the present invention implements SLA requirements that must be met before certain payments in the debit/credit system are made or refunded.

Certain embodiments may include multiple backbones, which may be partitioned to offer differentiated services. Such differentiated services might be based on geography, industry, or some other attribute or categorization. Each differentiated service would preferably have its own respective debits/credits/brokerage fees. In addition, unique debit/credit/brokerage fee structures may be applied to communications between customers that require traversing two or more of these partitioned backbones. For example, there may be a North American backbone and a Pan-European backbone. Different debit/credit/brokerage fees may be charged if a particular message crosses from one partitioned backbone to another.

Accordingly, the present invention is intended to embrace all such alternatives, modifications and variations as fall within the spirit and scope of the appended claims and equivalents.